

What Is Claimed Is:

- 1 1. A method for logging file system operations, comprising:
2 receiving a request to perform a file system operation;
3 making a call to an underlying file system to perform the file system
4 operation; and
5 logging the file system operation to a log within a log device to facilitate
6 recovery of the file system operation in the event of a system failure before the file
7 system operation is committed to non-volatile storage.
- 1 2. The method of claim 1, wherein logging the file system operation
2 involves storing an identifier for the file system operation to the log device.
- 1 3. The method of claim 1, further comprising periodically committing
2 the log to the underlying file system by:
3 freezing ongoing activity on a file system;
4 making a call to the underlying file system to flush memory buffers to non-
5 volatile storage, whereby outstanding file system operations are guaranteed to be
6 committed to non-volatile storage;
7 removing outstanding file system operations from the log; and
8 unfreezing the ongoing activity on the file system.
- 1 4. The method of claim 1, wherein upon a subsequent computer
2 system startup, the method further comprises:
3 examining the log within the log device;
4 replaying any file system operations from the log that have not been
5 committed to non-volatile storage.

1 5. The method of claim 1, further comprising checking for
2 dependencies between the file system operation and ongoing file system
3 operations; and
4 if dependencies are detected, ensuring that the file system operation and
5 the ongoing file system operations complete in an order that satisfies the
6 dependencies.

1 6. The method of claim 1,
2 wherein the request to perform the file system operation is received at a
3 primary server in a highly available system; and
4 wherein the log device includes a secondary server in the highly available
5 system that acts as a backup for the primary server.

1 7. The method of claim 1, further comprising:
2 associating the file system operation with a transaction identifier for a set
3 of related file system operations; and
4 wherein logging the file system operation involves storing the file system
5 operation with the transaction identifier to the log device.

1 8. The method of claim 1, wherein logging the file system operation
2 involves:
3 determining if the file system operation belongs to a subset of file system
4 operations that are subject to logging; and
5 if so, logging the file system operation.

1 9. The method of claim 8, wherein the subset of file system
2 operations are non-idempotent file system operations.

1 10. The method of claim 1, wherein the log device stores the file
2 system operation in volatile storage.

1 11. The method of claim 1, wherein the log device stores the file
2 system operation in non-volatile storage.

1 12. A computer-readable storage medium storing instructions that
2 when executed by a computer cause the computer to perform a method for logging
3 file system operations, the method comprising:
4 receiving a request to perform a file system operation;
5 making a call to an underlying file system to perform the file system
6 operation; and
7 logging the file system operation to a log within a log device to facilitate
8 recovery of the file system operation in the event of a system failure before the file
9 system operation is committed to non-volatile storage.

1 13. The computer-readable storage medium of claim 12, wherein
2 logging the file system operation involves storing an identifier for the file system
3 operation to the log device.

1 14. The computer-readable storage medium of claim 12, wherein the
2 method further comprises periodically committing the log to the underlying file
3 system by:
4 freezing ongoing activity on a file system;

5 making a call to the underlying file system to flush memory buffers to non-
6 volatile storage, whereby outstanding file system operations are guaranteed to be
7 committed to non-volatile storage;
8 removing outstanding file system operations from the log; and
9 unfreezing the ongoing activity on the file system.

1 15. The computer-readable storage medium of claim 12, wherein upon
2 a subsequent computer system startup, the method further comprises:
3 examining the log within the log device;
4 replaying any file system operations from the log that have not been
5 committed to non-volatile storage.

1 16. The computer-readable storage medium of claim 12, wherein the
2 method further comprises checking for dependencies between the file system
3 operation and ongoing file system operations; and
4 if dependencies are detected, ensuring that the file system operation and
5 the ongoing file system operations complete in an order that satisfies the
6 dependencies.

1 17. The computer-readable storage medium of claim 12,
2 wherein the request to perform the file system operation is received at a
3 primary server in a highly available system; and
4 wherein the log device includes a secondary server in the highly available
5 system that acts as a backup for the primary server.

1 18. The computer-readable storage medium of claim 12, wherein the
2 method further comprises:

3 associating the file system operation with a transaction identifier for a set
4 of related file system operations; and
5 wherein logging the file system operation involves storing the file system
6 operation with the transaction identifier to the log device.

1 19. The computer-readable storage medium of claim 12, wherein
2 logging the file system operation involves:
3 determining if the file system operation belongs to a subset of file system
4 operations that are subject to logging; and
5 if so, logging the file system operation.

1 20. The computer-readable storage medium of claim 19, wherein the
2 subset of file system operations are non-idempotent file system operations.

1 21. The computer-readable storage medium of claim 12, wherein the
2 log device stores the file system operation in volatile storage.

1 22. The computer-readable storage medium of claim 12, wherein the
2 log device stores the file system operation in non-volatile storage.

1 23. An apparatus that logs file system operations, comprising:
2 a receiving mechanism that is configured to receive a request to perform a
3 file system operation;
4 a calling mechanism that is configured to make a call to an underlying file
5 system to perform the file system operation; and
6 a logging mechanism that is configured to log the file system operation to
7 a log within a log device to facilitate recovery of the file system operation in the

8 event of a system failure before the file system operation is committed to non-
9 volatile storage.

1 24. The apparatus of claim 23, wherein the logging mechanism is
2 configured to store an identifier for the file system operation to the log device.

1 25. The apparatus of claim 23, wherein the logging mechanism is
2 configured to periodically:
3 freeze ongoing activity on a file system;
4 make a call to the underlying file system to flush memory buffers to non-
5 volatile storage, whereby outstanding file system operations are guaranteed to be
6 committed to non-volatile storage;
7 remove outstanding file system operations from the log; and to
8 unfreeze the ongoing activity on the file system.

1 26. The apparatus of claim 23, further comprising a recovery
2 mechanism that operates during system startup, wherein the recovery mechanism
3 is configured to:
4 examine the log within the log device; and to
5 replay any file system operations from the log that have not been
6 committed to non-volatile storage.

1 27. The apparatus of claim 23, further comprising a dependency
2 handler that is configured to:
3 check for dependencies between the file system operation and ongoing file
4 system operations; and to

5 ensure that the file system operation and the ongoing file system
6 operations complete in an order that satisfies dependencies if dependencies are
7 detected.

1 28. The apparatus of claim 23,
2 wherein the receiving mechanism is located within a primary server in a
3 highly available system; and
4 wherein the log device is located within a secondary server in the highly
5 available system that acts as a backup for the primary server.

1 29. The apparatus of claim 23, further comprising a transaction
2 mechanism that is configured to associate the file system operation with a
3 transaction identifier for a set of related file system operations; and
4 wherein the logging mechanism is configured to log the file system
5 operation with the transaction identifier to the log device.

1 30. The apparatus of claim 23, wherein the logging mechanism is
2 configured to:
3 determine if the file system operation belongs to a subset of file system
4 operations that are subject to logging; and to
5 log the file system operation if the file system operation belongs to the
6 subset of file system operations that are subject to logging.

1 31. The apparatus of claim 30, wherein the subset of file system
2 operations are non-idempotent file system operations.

1 32. The apparatus of claim 23, wherein the log device is configured to
2 store the file system operation in volatile storage.

1 33. The apparatus of claim 23, wherein the log device is configured to
2 store the file system operation in non-volatile storage.